A Linear Dual-Space Approach to 3D Surface Reconstruction from Occluding Contours using Algebraic Surfaces

Kongbin Kang	Jean-Philippe Tarel	Richard Fishman	David Cooper
Div. of Eng.	LIVIC (INRETS-LCPC)	Program in Visual Arts	Div. of Eng.
Brown University	13 Route de la Miniere	Brown University	Brown University
Providence, RI 02912	78000 Versailles, France	Providence RI 02912	Providence, RI, 02912

Abstract

We present a linear approach to the 3D reconstruction problem from occluding contours using algebraic surfaces. The problem of noise and missing data in the occluding contours extracted from the images leads us to this approach. Our approach is based first on the intensive use of the duality property between 3D points and tangent planes, and second on the algebraic representation of 3D surfaces by implicit polynomials of degree 2 and higher.

1 Introduction

Surface reconstruction from images has been intensively investigated over the past three decades. We focus here on the 3D reconstruction from the occluding contours in the image only. Developments of the early work [5] on this special kind of reconstruction, under known and small camera motion, are compared in [2]. Another kind of approach was proposed in [6, 3, 9, 7] where camera motion can be large between views, but where the occluding contour (also called apparent contour) is assumed cleanly extracted from the images. These techniques were extended to unknown camera motion, see for instance [1, 2, 7]. Many of these techniques make interesting, innovative effective use of the duality property between points and planes in homogeneous coordinates. Indeed, this duality is particularly useful to handle 3D object surface reconstruction from occluding contours. But it is only in [10] that optimally combining all the available measurements for better robustness to noise is taken into account and partially solved. Another practical problem also arises from missing data in the occluding contours.

On the way to our long-term goal to bring surface reconstruction from gray-level, occluding contour and surface albedo-discontinuity data approaches together in a unified Bayesian treatment, we present a new approach to 3D reconstruction from occluding contours which *has the advantage of being linear*, does not require point matching, is robust to noise and missing data on the occluding contours, and is model-based, thus producing accurate reconstruction that explicitly captures important geometry and is of low computational-cost. Our approach is based first on the intensive use of the duality property, and second on the algebraic representation of 3D surfaces by implicit polynomials. It is the conjunction of these two concepts that leads to the proposed linear solution for tackling contour segmentation problems. Note, [6] and others [3], introduced an approach to computing quadric surface patches that is linear in the coefficients of quadratic curves that are fit to occluding contours in the images used. A very important practical consideration is that the quadratic curve coefficients may not be stable in general unless a long patch of data is used, which may not be feasible, particularly when data are noisy or have missing parts. Our quadric surface-patch linear-estimation is based on tangent-line fits to the occluding contours in all the image, and is therefore less vulnerable to the preceding perturbations. Also our approach is extensible to algebraic surfaces of higher degree.

The content of this paper is as follows: First the properties of the envelope of a 2D curve (Sec. 2) and of a 3D surface (Sec. 3) are reviewed. Then, the algebraic surface representation is described [12] (Sec. 4). Sec. 5 is dedicated to the particular case of the quadrics which was investigated in a different way ([6, 3]). In Sec. 6, we explain how to make good use of the duality property and algebraic surfaces of higher degree for reconstructing a 3D surface from its occluding contours in a particularly simple linear way. Indeed, it is well known (see Fig. 1) that, given the 3×4 calibration matrix M of a camera, the tangent plane to the surface $\mathcal{S}$ to be reconstructed can be estimated completely from the tangent line in the image of the camera.

However, the 3D position of contact point X between tangent plane P and $\mathcal{S}$ is unknown. We only know its image $U = MX$ in homogeneous coordinates. This means that the 3D reconstruction from occluding contours can be seen

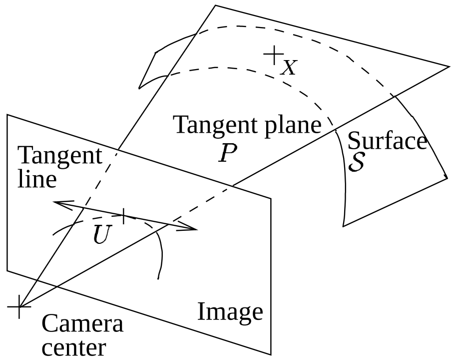

Figure 1. Tangent plane of a 3D point X, its image U, and the tangent line to the occluding contour in the image.

as the search for the surface $\mathcal{S}$ given a data set of tangent planes, i.e., given a sampling of the envelope of $\mathcal{S}$. In this paper, we develop fitting in the space of planes $\mathcal{E}$, using the algebraic surface representation (see [11])

Sec. 7 describes our approach to partitioning 3D scene-space into an array of cubes and estimating a 3D quadric surface patch for the object surface lying in a cube. Sec. 8 describes the use of quadric patches, found in Sec. 7, in estimating a single higher degree algebraic surface to represent all or a subset of the object surface patches lying in cubes. Experiments illustrating these ideas are presented in Sec. 9. In Sec. 10 and 11, we touch on concepts for improving representation accuracy and stability when estimating high degree algebraic surfaces in the dual space directly from the data.

2 Envelope of a Parameterized Curve and Duality

By definition, the envelope of a planar curve is the entire set of tangent lines. Given a planar curve $X(\lambda) = \begin{bmatrix} x(\lambda) \\ y(\lambda) \end{bmatrix}$, parametrized by λ, the tangent vector at point $X(\lambda)$ is $\frac{dX}{d\lambda}(\lambda)$. Thus the equation of the tangent line at this point is:

$$\left[\left(\frac{d}{d\lambda}X(\lambda)\right)^{\perp}\right]^t (X' - X(\lambda)) = 0 \qquad (1)$$

where λ is fixed, $\left(\frac{dX}{d\lambda}(\lambda)\right)^{\perp}$ is a vector normal to the curve at $X(\lambda)$ and X' is the variable 2D point on the tangent line. Eq. (1) is defined up to a scale factor μ. Therefore the envelope can be represented as the complete set of 3D vectors:

$$L(\lambda,\mu) = \mu \begin{bmatrix} \left(\frac{d}{d\lambda}X(\lambda)\right)^{\perp} \\ -\left[\left(\frac{d}{d\lambda}X(\lambda)\right)^{\perp}\right]^t X(\lambda) \end{bmatrix} \qquad (2)$$

The envelope can be represented as a 2D surface parameterized by λ and μ in a 3D space, called the space of lines. Due to μ, the envelope is a generalized conic surface, i.e., any ray going through the origin and another point of the envelope is included in the envelope. To introduce the concept of duality, we notice from Eq. (2) that:

$$L^t(\lambda,\mu) \begin{bmatrix} \frac{d}{d\lambda}X(\lambda) \\ 0 \end{bmatrix} = 0 \qquad (3)$$

Moreover, by definition the point $X(\lambda)$ is always on the tangent line $L(\lambda,\mu)$, i.e., $L^t(\lambda,\mu)\begin{bmatrix} X(\lambda) \\ 1 \end{bmatrix} = 0$. By differentiation, we deduce:

$$\frac{\partial}{\partial\lambda}L^t(\lambda,\mu)\begin{bmatrix} X(\lambda) \\ 1 \end{bmatrix} + L^t(\lambda,\mu)\begin{bmatrix} \frac{d}{d\lambda}X(\lambda) \\ 0 \end{bmatrix} = 0$$

The last term cancels using Eq. (3), and therefore:

$$\frac{\partial}{\partial\lambda}L^t(\lambda,\mu)\begin{bmatrix} X(\lambda) \\ 1 \end{bmatrix} = 0 \qquad (4)$$

¿From Eq. (2), we deduce:

$$\frac{\partial L^t}{\partial\mu}(\lambda,\mu)\begin{bmatrix} X(\lambda) \\ 1 \end{bmatrix} = 0 \qquad (5)$$

The envelope being a generalized conic surface, any tangent plane is going through the origin and thus is completely characterized by its normal. The last two equations mean that the normal of the tangent plane of the envelope at $L(\lambda,\mu)$ (i.e., $L(\lambda,\mu)$ is the tangent line of the curve at $X(\lambda)$) is oriented by $\begin{bmatrix} X(\lambda) \\ 1 \end{bmatrix}$. This property is the so called duality property between points and lines [8]. It is the first property we are taking advantage of in our linear approach to 3D reconstruction from occluding contours.

3 Envelope of a Parameterized Surface

The duality property extends to 3D surfaces. Let $X(\lambda,\nu)$ be a parameterized surface, and $N(\lambda,\nu)$ the normal vector at the 3D point $X(\lambda,\nu) = \begin{bmatrix} x(\lambda,\nu) \\ y(\lambda,\nu) \\ z(\lambda,\nu) \end{bmatrix}$. The equation of the tangent plane at $X(\lambda,\nu)$ is $N^t(\lambda,\nu)(X' - X(\lambda,\nu)) = 0$, where X' is the variable point on the tangent plane. The envelope is the hyper-surface:

$$P(\lambda,\mu,\nu) = \mu \begin{bmatrix} N(\lambda,\nu) \\ -N^t(\lambda,\nu)X(\lambda,\nu) \end{bmatrix}$$

in the 4D space (p,q,r,s), the space of planes. Similarly as in the previous section, we can prove that the normal of the tangent plane of the envelope at $P(\lambda,\mu,\nu)$ (i.e., $P(\lambda,\mu,\nu)$ is the tangent plane of the original surface at $X(\lambda,\nu)$) is $\begin{bmatrix} X(\lambda,\nu) \\ 1 \end{bmatrix}$.

4 Envelope of an Algebraic Surface

To get rid of the parameterization that generally does not lead to a simple fitting technique, we use the algebraic representation of surface. Formally, an algebraic 3D surface is specified by a 3D Implicit Polynomial (IP) surface of degree n' given by the following equation:

$$f_{n'}(x,y,z) = \sum_{0\le i+j+k\le n'} a_{ijk}x^i y^j z^k = 0$$

The envelope $\mathcal{E}$ of any algebraic surface $\mathcal{S}$ is an algebraic surface as demonstrated by elimination theory [4]. The algebraic 4D envelope is specified by a 4D *homogeneous* IP of degree n,where generally $n \neq n'$, given by the following equation:

$$g_n(p,q,r,s) = \sum_{0\le i+j+k+l\le n} a_{ijkl}p^i q^j r^k s^l = 0$$

in the space of planes. The polynomial must be homogeneous since the envelope is a generalized conic surface.

Elimination theory can be used to derive analytically the polynomial equation of envelope $\mathcal{E}$ given the IP equation of original surface $\mathcal{S}$ [4]. Second degree polynomials lead to simple equations as shown in [3] and in the next section.

Rather than using pure algebraic techniques, we use the duality property to sample $\mathcal{S}$, given $\mathcal{E}$. Given any point $P(p,q,r,s)$ on envelope $\mathcal{E}$, the normal direction $N(P)$ to $\mathcal{E}$ at P is obtained directly by taking the gradient of $g_n(p,q,r,s)$ with respect to all the coordinates:

$$\nabla g_n(p,q,r,s) = \begin{bmatrix} \frac{\partial}{\partial p} \\ \frac{\partial}{\partial q} \\ \frac{\partial}{\partial r} \\ \frac{\partial}{\partial s} \end{bmatrix} g_n(p,q,r,s)$$

Duality implies that $\nabla g_n(P)$ is the homogeneous coordinates of the contact point between plane P and $\mathcal{S}$. In practice, polynomial $g_n(p,q,r,s)$ is represented by the coefficient vector $(a_{ijkl})_{0\le i,j,k,l;i+j+k+l=n}$ which has dimension $p = \frac{1}{6}(n+1)(n+2)(n+3)$, i.e. number of coefficients:

$$g_n(p,q,r,s) = Y^t A \quad (6)$$

where A is the ordered coefficient vector and Y the monomial vector.

5 Quadrics

An algebraic 3D surface of degree 2 is a quadric. The equation of a quadric can be rewritten as:

$$\begin{aligned} f_2(X) &= X^t B X + 2C^t X + d \\ &= \begin{bmatrix} X \\ 1 \end{bmatrix}^t \underbrace{\begin{bmatrix} B & C \\ C^t & d \end{bmatrix}}_{D} \begin{bmatrix} X \\ 1 \end{bmatrix} = 0 \end{aligned}$$

where B is a 3×3 matrix, C a 3D vector, and d a scalar. The gradient of $f_2(X)$ is $2BX+2C$. Therefore, the tangent plane $P(X)$ at X is $P(X) = 2\begin{bmatrix} BX+C \\ -(BX+C)^t X \end{bmatrix}$. One can check that for every X:

$$P^t(X)\begin{bmatrix} B & C \\ C^t & d \end{bmatrix}^{-1} P(X) = 0$$

This means that envelope $\mathcal{E}$ of a quadric is a second degree homogeneous polynomial on the 4D space of planes. Its matrix is simply the inverse of the matrix D of the quadric in homogeneous coordinates.

When the envelope matrix D^{-1} is known, for any tangent plane P, the contact point $X(P)$ is obtained by the following formula

$$k\begin{bmatrix} X(P) \\ 1 \end{bmatrix} = D^{-1}P \quad (7)$$

where k is an arbitrary constant. This formula allows us to sample the contour generator (i.e. the curve in the 3D surface which generate the occluding contour) as explained next.

6 Fitting in the Space of Planes

At an image point U of the contour generator of the surface $\mathcal{S}$ to be reconstructed, the tangent line is computed, and using the camera calibration a normalized tangent plane $P(U)$ of $\mathcal{S}$ is estimated. As defined previously, all the 4D points $\mu P(U)$ are points on envelope $\mathcal{E}$.

This means that by working in the space of planes rather than staying in the original 3D space as in the classical approaches, we set the problem as fitting a 4D hyper-surface, i.e. the envelope $\mathcal{E}$, on 4D data points which represent 3D planes in the original 3D space. When the fitting of $\mathcal{E}$ is performed, using the duality property, any normal to $\mathcal{E}$ gives up to a scale factor a point on $\mathcal{S}$. Indeed, $\mathcal{S}$ has been implicitly reconstructed by fitting its envelope $\mathcal{E}$ in the space of planes. This is how to compute point estimates on $\mathcal{S}$ from the algebraic surface fit in the dual space.

We choose to use the algebraic representation of the 4D surfaces since it leads to linear fitting techniques. The classical and simplest way to fit an algebraic surface to data is to minimize the algebraic distance over the set of given data points $P_j = (p_j, q_j, r_j, s_j)$, $1 \le j \le m$, that is

$$e_{algebraic} = \sum_{j=1}^{m}(g_n(P_j))^2 = A^t \underbrace{\left(\sum_{1\le j\le m} Y_j Y_j^t\right)}_{V} A \quad (8)$$

by using the vector representation of g_n as in (6). The symmetric matrix V is the so-called scatter matrix.

To avoid the trivial zero solution in the minimization of (8), a constraint such as $||A||^2 = 1$ is imposed which modifies the problem to

$$\min_A \left(A^t \left(\sum_{1 \le j \le m} Y_j Y_j^t \right) A + \gamma (A^t A - 1) \right) \quad (9)$$

with the introduction of Lagrangian multiplier γ. The solution to (9) is given by the unit eigenvector A associated with γ_{min}, the smallest eigenvalue of $VA = \gamma A$. Consequently, the classical least-squares fitting algorithm consists of computing the monomial scatter matrix V from a set of data points, and then finding the unit eigenvector of V associated with its smallest eigenvalue. In [11] fitting techniques more stable under noise and missing data are described.

7 Practical Estimation of a Complex Surface Based on the Union of Estimated Quadric Patches

For simplicity of presentation, we illustrate the approach in Fig. 2 in 2D where the scene space is a plane which is decomposed into squares, and the object boundary is a curve in the plane. Four rays 1, 2, 3, 4 are shown. Each is tangent to the object surface at a point on its occluding contour and each lies in a different image from the set of images I_1, I_2, I_3 and I_4. The approach can be thought of as space-carving at a coarse level and model-based surface fitting at a fine level.

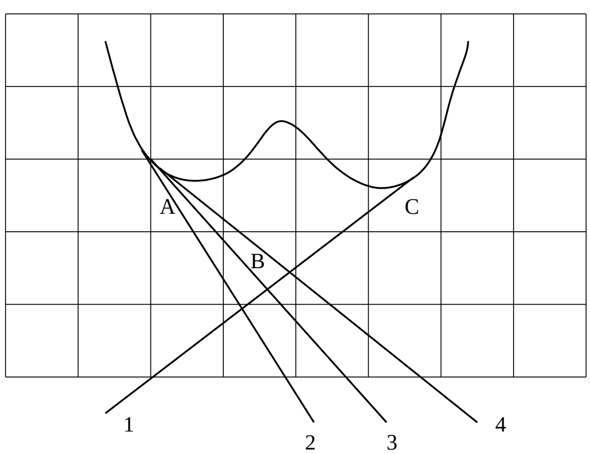

Figure 2. 2D illustration of 3D decomposition of scene-space into cubes, and object boundary in this space.

To estimate a large complex 3D surface, we a priori partition the scene space into an array of cubes, and estimate the 3D surface patch lying within each cube. The advantage of this approach is the simplicity of obtaining estimated surface patches that cover the entire portion of the object-surface that is seen in the images, and the necessity of only a relatively small amount of data for estimating a quadric patch.

Consider cube A in Fig. 2. Project its boundaries into images I_1, I_2, I_3 and I_4. Within the projected convex boundary regions, occluding contours are seen in images I_2, I_3, I_4 only. Estimate the tangent lines to the occluding contours seen. Take the 3D planes passing through these lines and the associated camera center of projection. These planes are the P_j's in Sec. 6. Following Sec. 5, the solution is obtained which specifies the estimated quadric surface from the data sets observed in the images. First the fitting error can be used to check that the quadratic patch is consistent with the P_j's. Moreover, using (7), we can easily compute the 3D points on the contour generators which correspond to the edge tangent lines on the occluding contours. Now check to see whether the 3D points on contour generators are mostly within the cube A. If yes, then we have an estimated 3D surface for that portion of the surface that lies within cube A. If no, then the occluding contours used are not of a portion of the true 3D surface that lies within cube A.

Now consider cube B. Project 3D cube boundary into images I_1, I_2, I_3, I_4. There are occluding contours in each of the four projected boundaries. Try estimating a quadric patch using all four images. First the fitting error becomes larger. And even if some portion of the determined global quadric surface intersects cube B, most of the contour generators of the estimated surface patch will not be contained in cube B. We can then consider a subset of the occluding contours, e.g., the three in images I_2, I_3, I_4. Again, the estimated surface will either not intersect cube B or not have an intersection which is consistent with occluding contours used in I_2, I_3, I_4.

The *check process is computationally cheap here*, since the computation incurred is only to get the contour generator points, and the homogeneous coordinates of 3D contour generator points can be obtained by only a multiplication between the estimated surface matrix and homogeneous coordinates of the tangent plane. Robust fitting techniques can be used to tackle the presence of outliers and badly segmented occluding contours.

8 Estimating a High Degree Algebraic Surface for a Complex Object from Two or More Locally Estimated Quadric Patches

Consider the object in Fig. 3. This is a deformation of a torus. The general shape can be captured by a single $4th$ degree algebraic surface and a good representation can probably be had with a single $6th$ degree surface. Alternatively, the surface can probably be simply decomposed into

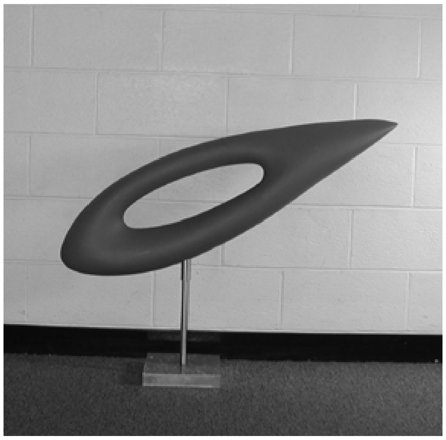

Figure 3. The sculpture used for the experiment

two surfaces, each connected and each well represented by a single $4th$ degree algebraic surface. The concept of interest here is that we want to capture geometry useful for representing and describing a 3D surface, and that means structure more complicated than planar triangular patches or quadric patches. Ridges, deformed toroids, deformed cylinders, etc., are more complex free-form primitives of interest, and $4th$ or $6th$ degree algebraic surfaces are good models for the purpose. One approach toward this end is fitting a higher degree algebraic surface to the quadric patches estimated in all or a subset of the surface-occupied scene-space cubes. This can be accomplished in either of two ways. One way is to estimate the coefficients of the high degree algebraic surface directly from the coefficients and their probability distributions for the quadric patches. This can be done, but it is involved. We take a $2nd$ simpler approach which is to take a large sample of points from the portion of each estimated quadric patch that lies within a scene-space cube and is close to the contour generators there associated with the occluding contours in images used. The high degree algebraic surface is then fit to this simulated data. We know which quadric patches to use from the quadric patch estimation approach discussed in Sec. 7. This approach is illustrated in Sec. 9.

9 Experiments with Quadric Patches

Fig. 3 is the image of the $3D$ object of interest. Our surface reconstruction algorithm assume a metrically calibrated sequence of images. The camera is calibrated by having the same planar calibration checkerboard pattern in each image used. This pattern is simply printed on paper and glued on a cardboard box which does not have a perfectly flat surface. Using the same planar pattern in all images severely limits the usable range of camera viewing positions, which is the reason we have reconstructed only the portion of object-surface shown in Figs. 5b and 6. To function more practically in a sculpture, archaeology or other environment so that the camera can use a much large range of viewing positions, we will either distribute a number of such planar calibration patterns in the work space, or use a single calibration cube with a pattern on each surface, or use an internally calibrated camera and a number of easily identifiable circular discs or line intersections distributed in the vicinity of the object. In this experiment, fourteen images were taken. Among them, four images - leftmost, rightmost, topmost and bottommost - are shown in Fig. 4.

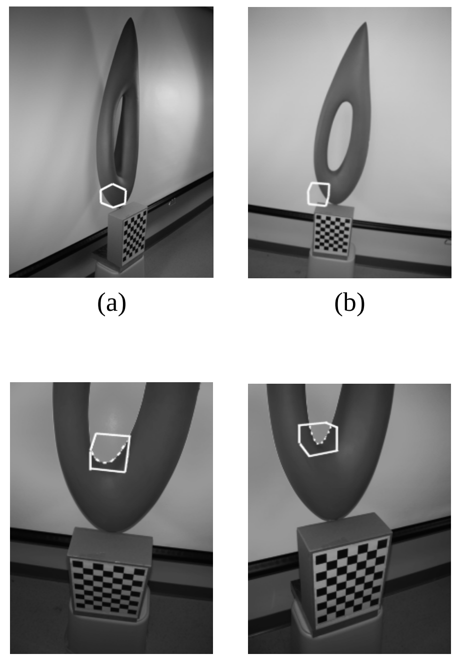

(a) (b)

(c) (d)

Figure 4. Four out of fourteen images used for reconstructing are shown here. These are four extreme right, left, down, up views. The white polygon is the convex hull of the projected cube boundary. The dotted white curves are the occluding contours of the reconstructed surface seen in two of the cubes.

From the 60 cubes considered, quadric patches were estimated in the 10 cubes that contained contour generators. In the Fig. 4, only two of these boxes (one is in 4a-b and the other is in 4c-d) are shown: the white convex polygon in each image is the projected boundary of the cube being investigated. The dotted white curves shown in the images are the occluding contour in the view seen of the estimated 3D quadric surface. Although all cubes have the same size

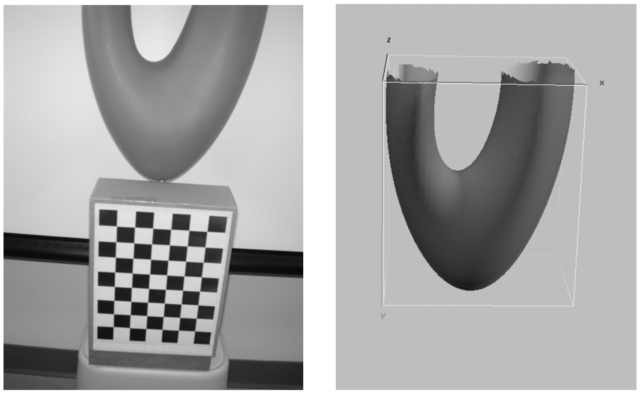

(a) (b)

Figure 5. (a) The bottom portion of the real sculpture. (b) The $4th$ degree algebraic surface reconstructed for the bottom portion

in 3D, they show not the same in the images because the distances of the camera to the sculpture change from view to view.

Fig. 5b is a $4th$ degree algebraic surface fit to data points sampled from the 10 quadric surface patches, one in each cube of interest. Note that the $4th$ degree surface captures the shape of the original sculpture Fig. 5a well in the vicinities of the 10 cubic regions of space of interest. This is further illustrated in Fig. 6 where the white points are samples from the ten estimated quadric patches and were used in estimating the $4th$ degree patch. Only those points used that lie slightly outside the $4th$ degree algebraic surface are seen. The others are slightly insider the surface. An interesting observation is that this $4th$ degree algebraic surface extrapolates the data in a meaningful way. This is illustrated in Fig. 7 where the entire zero set of the fitted polynomial is shown. Although data from only the lower 3rd of the true surface is used in estimating the $4th$ degree polynomial, the combination of the data and restriction of the reconstructed surface to a $4th$ degree polynomial results in the upper portion of the reconstruction capturing the essence of the upper portion of the object.

10 On Using all of the Information in the Images in Fitting Algebraic Surfaces in the Dual Space Directly

It is important to use all the information that is available. The fitting solution given in (8) minimizes algebraic distance in the tangent-plane space. This is not the most accurate distance measure to use with noisy data. We therefore propose introducing two additional terms in the objective function to be minimized. These additional terms constitute the squared distance between the occluding contours of the surface estimate and data points on the occluding contours in the images of the object.

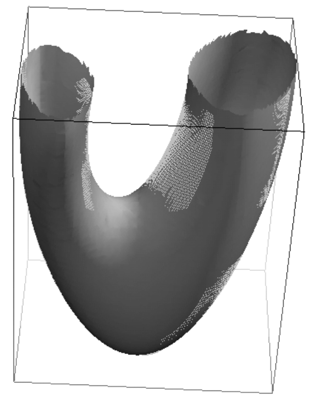

Figure 6. $4th$ degree algebraic surface fit to data points sampled from the two quadric surface patches. The white points in the surface are the data points sampled from two quadric patches.

Indeed, we know that 4D normal $N(P(U))$ to $\mathcal{E}$ at $P(U)$ equals, up to a scale factor, the homogeneous coordinates X of the contact point of $\mathcal{S}$ and $P(U)$ in the original 3D space. Therefore, we deduce that $MN(P(U))$ must be equal to U up to a scale factor, where M is the perspective projection matrix of the camera. This gives two constraints on the gradient of the algebraic representation of $\mathcal{E}$ at $P(U)$. These two constraints can be written in a linear way with respect to the parameters of the algebraic representation of $\mathcal{E}$. Let M_1, M_2, and M_3, denote the lines of M. These constraints add two terms to (8) to yield

$$e_{recons} = \sum_{j=1}^{m} ((g_n(P_j))^2 + \alpha(((M_1 - u_j M_3)\nabla g_n(P_j))^2 + ((M_2 - v_j M_3)\nabla g_n(P_j))^2)) \qquad (10)$$

where $U_j = (u_j, v_j)$ is the image point where the tangent plane P_j is estimated and α is the relative weights on the gradient with respect to the g^2 term. Following Bayesian theory, a good choice for α is the one that weights equally the average errors on the image tangent lines and image points.

11 Improving the Stability of Fitting Algebraic Surfaces in the Dual Space

Stabilization of estimated algebraic curves and surfaces and their coefficients generally requires the use of ridge re-

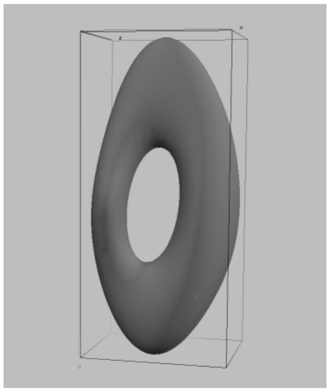

Figure 7. The restructed surface using $4th$-degree algebraic surface and data over portions of the lower $1/3$ of the object only (see Fig. 6). Although this reconstruction does not very accurately conform to the original shape over its upper half, the reconstruction still captures the object's main features.

gression [11]. This is especially true when fitting high degree polynomials or when only a small part of the data set is available. We are experimenting with stabilization in our current work. To stabilize our estimates, we have to decide what kind of a prior information we want to introduce in the fitting, i.e., towards which shape we want to bias the result of the fitting. For higher degree, if the estimated S should be close to a quadric with matrix D (eventually estimated as in Sec. 9), its envelope should be close to a second degree algebraic surface with matrix D^{-1} as explained in Sec. 5. Preliminary experiments are promising.

12 Conclusion

A linear low computational-cost algorithm and theory were introduced for estimating 3D quadric surfaces from estimated tangent lines to the occluding contours of an object in 3 or more images. In general, images taken from 3 positions are necessary for achieving this reconstruction [6]. Our approach is quadric surface fitting in the 4th dimensional homogeneous dual space consisting of tangent planes.

To capture useful descriptive structure, we introduced the estimation of a complex-object surface by a union of locally estimated 3D quadric patches followed by fitting a single higher degree algebraic surface to all or many of these quadric patches, and illustrated this approach. We introduced, but did not describe, an implementation of an alternative concept which is to fit a single higher degree algebraic surface in the dual space of 3D planes, which then uniquely determines a higher degree 3D algebraic surface representing the 3D surface seen in the images. 3D points on this algebraic surface can be simply computed.

The approaches described in this paper open up a number of avenues for further research. Among these is new approaches for automatic tracking of occluding contours in a sequence of video frames. This is necessary in order to have a practical accurate surface reconstruction and significant-geometry extraction system. Another is comparison of 3D surface reconstruction accuracy of this linear approach versus a maximum accuracy nonlinear approach directly in the image spaces. This comparison is realizable due to the surface models we use and applying large sample theory and perturbation analysis. A third avenue is most useful representations for complex surfaces and their significant geometry.

Acknowledgement

This work was partially supported by NSF Grants BCS-9980091 and IIS-9802392.

References

[1] K. Åström and F. Kahl. Motion estimation in image sequences using the deformation of apparent contours. *Trans. PAMI*, 21(2):114–127, 1999.

[2] G. Cross, A. Fitzgibbon, and A. Zisserman. Parallax geometry of smooth surfaces in multiple views. In *Proc. ICCV*, pages 323–329, 1999.

[3] G. Cross and A. Zisserman. Quadric surface reconstruction from dual-space geometry. In *Proc. ICCV*, pages 25–31, 1998.

[4] D. Forsyth. Recognizing algebraic surfaces from their outlines. In *Proc. ICCV*, pages 476–480, 1993.

[5] P. Giblin and R. Weiss. Reconstructions of surfaces from profiles. In *Proc. ICCV*, pages 136–144, 1987.

[6] S. Ma and L. Li. Ellipsoid reconstruction from three perspective views. In *Proc. ICPR*, page A8M.5, 1996.

[7] P. Mendonca, K. Wong, and R. Cipolla. Camera pose estimation and reconstruction from image profiles under circular motion. In *Proc. ECCV*, pages II: 864–877, 2000.

[8] D. Renaudie, D. Kriegman, and J. Ponce. Duals, invariants, and the recognition of smooth objects from their occluding contour. *Proc. ECCV*, 2000.

[9] S. Sullivan and J. Ponce. Automatic model construction and pose estimation from photographs using triangular splines. *Trans. PAMI*, 20(10):1091–1096, 1998.

[10] R. Szeliski and R. Weiss. Robust shape recovery from occluding contours using a linear smoother. *IJCV*, 28(1):27–44, 1998.

[11] T. Tasdizen, J.-P. Tarel, and D. B. Cooper. Improving the stability of algebraic curves for applications. *Trans. IP*, 9(3):405–416, 2000.

[12] G. Taubin. Estimation of planar curves, surfaces and nonplanar space curves defined by implicit equations, with applications to edge and range image segmentation. *Trans. PAMI*, 13(11):1115–1138, 1991.